UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/575,839	05/22/2000	Matt Ayers	52224/296056	2761

23370   7590   09/17/2008
JOHN S. PRATT, ESQ
KILPATRICK STOCKTON, LLP
1100 PEACHTREE STREET
ATLANTA, GA 30309

EXAMINER
NGUYEN, THANH T

ART UNIT	PAPER NUMBER
2144	

MAIL DATE	DELIVERY MODE
09/17/2008	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

1   RECORD OF ORAL HEARING
2
3   UNITED STATES PATENT AND TRADEMARK OFFICE
4   _____
5
6   BEFORE THE BOARD OF PATENT APPEALS
7   AND INTERFERENCES
8   _____
9
10   *Ex parte* MATT AYERS, BEN BLACK, CHRIS BROWN, JOHN
11   CARLSON, DAN COHN, SCOTT LAIRD, JON MILLER, STEPHEN
12   RAMSEY, OPHIR RONEN, PAUL SCHACTER and OSCAR
13   STIFFELMAN
14   _____
15
16   Appeal 2008-2560
17   Application 09/575,839
18   Technology Center 2100
19   _____
20
21   Oral Hearing Held: August 13, 2008
22   _____
23
24
25   Before JEAN R. HOMERE, ST. JOHN COURTENAY III, and STEPHEN
26   C. SIU, Administrative Patent Judges.
27
28   ON BEHALF OF THE APPELLANTS:
29
30   BRENDA O. HOLMES, ESQ.
31   Kilpatrick Stockton, LLP
32   1100 Peachtree Street
33   Atlanta, Georgia 30309
34
35
36   The above-entitled matter came on for hearing on Wednesday, August
37   13, 2008, commencing at 10:15 a.m., at the U.S. Patent and Trademark

1   Office, 600 Dulany Street, 9th Floor, Alexandria, Virginia, before Jennifer

2   M. O'Connor, Notary Public.

3

4   THE CLERK: Calendar number 19, Mrs. Holmes.

5   JUDGE HOMERE: Good morning, Counselor.

6   MS. HOLMES: Good morning.

7   JUDGE HOMERE: You have 20 minutes. Please feel free to begin

8   any time you're ready.

9   MS. HOLMES: Okay, thank you. We're here this morning to discuss

10   a case that is related to routing across a network and it does this in a -- the

11   routing scheme that this particular invention uses is a combination of using

12   cost measurements that measure the operational characteristics of a network,

13   things like latency or loss, those types of things. But it does it in

14   combination with using physical proximity information related to clients.

15   So just to set up the scenarios for you. If you have multiple clients

16   and then you would have multiple content servers so that there would be

17   multiple servers that provide the content requested by any one of these

18   clients. The invention is directed to selecting a particular server that would

19   make the most sense to provide the content in the most efficient way.

20   JUDGE HOMERE: Okay, what does that mean exactly? Are we

21   saying that -- you said something about proximity?

22   MS. HOLMES: Yes.

23   JUDGE HOMERE: Are you saying that -- I guess one would pick the

24   client, that would minimize our cost, the cost for the communication?

1   MS. HOLMES: It could be that, but given the fact that one could not

2   have cost or measurement data between all endpoints, between all potential

3   clients and all potential content servers --

4   JUDGE HOMERE: I understand that. Before --

5   MS. HOLMES: You have to select what you're going to measure and

6   how you're going to use those measurements.

7   JUDGE HOMERE: Before we even get to that, I want to focus on the

8   measurement aspect for a minute.

9   MS. HOLMES: Okay.

10   JUDGE HOMERE: So from what you're saying, if I understand well,

11   that if you have the proximity information, then that would amount to having

12   the cost information as well?

13   MS. HOLMES: Respectfully, no. The proximity information is

14   proximity between the clients, not between a particular client and any one or

15   more of the content servers. So what we see in the prior art is in particular

16   with the *Colby* reference.

17   JUDGE HOMERE: Before we go into the prior art, in your

18   specification you say the specification describes -- defines the cost

19   measurement to include proximity information.

20   MS. HOLMES: It could include proximity information, but it is not

21   limited to proximity information.

22   JUDGE HOMERE: In other words, if you have proximity

23   information, you should be able to derive cost information?

24   MS. HOLMES: You could use proximity information to make a cost

25   determination, yes.

3

1   JUDGE HOMERE: Therefore, if you had the proximity information,

2   you had the cost information?

3   MS. HOLMES: Not necessarily. That is -- that is one -- one way of

4   looking at this, but the invention is not --

5   JUDGE HOMERE: Hold on a second. You either have it or you

6   don't. Your specification says that cost information includes proximity

7   information; is that correct?

8   MS. HOLMES: Yes.

9   JUDGE HOMERE: Therefore, if you have proximity information,

10   you should be able to derive cost information; isn't that correct?

11   MS. HOLMES: Yes, you could.

12   JUDGE HOMERE: Therefore, proximity information equals -- or

13   would lead one of ordinary skill in the art to cost information, right?

14   MS. HOLMES: In one embodiment, yes.

15   JUDGE HOMERE: Okay, so we're clear on that. Okay, go ahead.

16   MS. HOLMES: Yes, but the invention is not limited to simply the use

17   of proximity information. And in fact, the specification does go on to say

18   that one of the challenges with the existing art, that it's only based on

19   proximity information. Simply picking two endpoints, a client and a content

20   server, based on proximity information does not always provide the best

21   performance, because simply using proximity information doesn't account

22   for other things that impact the performance.

23   And so the specification goes on to describe that there are

24   measurement techniques that can be used to measure the performance of a

25   content server relative to a particular client and to use that information for

1   routing decisions or -- not perhaps routing decisions, but for content server

2   selection decisions.

3   And so the specification goes on to describe that there are particular

4   types of scanners that can be used that are associated with a particular

5   content server, and that can measure, for example, delay between certain

6   handshaking signals and can use that information to then say, this is a cost

7   associated with using this server with this particular client.

8   JUDGE HOMERE: This is great. I mean, we appreciate all this

9   background information, but let's focus -- let's try to remain focused here on

10   what the claim requires.

11   MS. HOLMES: Yes.

12   JUDGE HOMERE: What the claim recites, and the claim invention

13   recites cost information, cost measurement?

14   MS. HOLMES: Yes.

15   JUDGE HOMERE: Yeah, okay.

16   MS. HOLMES: And in particular, it recites cost measurements that

17   measure the operational characteristics of the network. I'm reciting from

18   claim 1.

19   JUDGE HOMERE: Okay. Okay. So how is that different from -- I

20   mean, earlier, just earlier we just went on an exchange and we agreed that if

21   you have proximity information, one of ordinary skill in the art would be

22   able to derive cost information. We know that the prior art teaches

23   proximity information, therefore -- based on our earlier exchange, therefore,

24   it teaches cost information. So what is lacking here?

25   MS. HOLMES: Respectfully, I believe that the claim language that

26   says cost measurements that measure operational characteristics direct the

1   claim language to the situation not related to simply proximity information,

2   but to the other types of cost measurements that are described in the

3   specification, such as --

4   JUDGE HOMERE: So you're saying that this claim language

5   excludes proximity information?

6   MS. HOLMES: I believe that it does, yes.

7   JUDGE SIU: Do you know -- where -- can you point to us where in

8   the specification there's a definition for cost measurements that excludes

9   proximity information?

10   MS. HOLMES: I will see if I can find that. I know that there are

11   several portions in the specification that talk about the limitations of using

12   simply proximity information between the client and the content server. I

13   don't know that it is phrased as stated, that cost excludes proximity.

14   JUDGE SIU: For example, on page four, lines 12 and 13, it says,

15   another example would be to use geographical location of the content server

16   as a performance criteria when they're talking about the measurements.

17   MS. HOLMES: Yes.

18   JUDGE SIU: So would that imply that geographical location or

19   proximity information would be part of that cost measurement definition?

20   MS. HOLMES: I believe that cost measurement could include

21   proximity information. I certainly don't disagree with your reading of the

22   specification there. I think that the language that we use to try to relate the

23   cost measurement to the operational characteristics was to focus on more of

24   the latency and delay types of characteristics as opposed to proximity.

25   JUDGE SIU: However, the way it's worded here, it seems like

26   proximity is included in there. If I'm reading a claim, and I'm reading this

6

1   portion of the specification, I would think that somebody might assume that

2   proximity information is included in cost measurements.

3   MS. HOLMES: I understand your comments and I guess, at least at

4   this point, that particular argument was one we didn't have an opportunity to

5   address with the examiner because it wasn't posed to us. But clearly we

6   were trying to focus on things other than proximity, because we have

7   consistently stated that proximity is not what we are trying to pursue.

8   JUDGE HOMERE: So are you saying that the focus here is on

9   operation or the measurement of operational characteristics of the network,

10   so by saying -- by reciting the operational -- measuring the operation

11   calculations of the network, that would not include proximity information?

12   MS. HOLMES: I believe that that is -- that is a fair interpretation of

13   the language.

14   JUDGE HOMERE: I guess the follow-up question would be, is there

15   anywhere in the specification where you say that operational information

16   includes delay, things of that nature?

17   MS. HOLMES: I'm sorry?

18   JUDGE HOMERE: Is there anywhere in the specification that give

19   some kind of guidance as to what operational characteristics are of the

20   network?

21   MS. HOLMES: There is a discussion that operational characteristics

22   would include network performance and I'm looking at page eight, line 29.

23   There is a discussion in that paragraph that begins at the bottom of page

24   eight and goes over to page nine that does talk about various types of cost.

25   JUDGE HOMERE: Performance is not necessarily synonymous with

26   -- okay. All right, go ahead.

1   MS. HOLMES: So that is certainly -- one aspect of the invention is

2 that the selection of a particular content server for a particular client is based

3 on this information that has been collected about network performance in the

4 past. If that information is available, it can then be used to select the content

5 server because you have relevant information for these two endpoints.

6   However, if that information is not available, then the claim goes on

7 to say that what we're going to do is we are going to use information that is

8 available from a different -- a distinct client and use that information to help

9 select the particular content server. And in the claim, we talk about the fact

10 that that second network client is physically proximate to the first network

11 client, and that is an area that we did focus on during prosecution with the

12 examiner.

13   JUDGE SIU: And how about since geographic location -- proximity

14 is one of the many factors you look at, so would it not have been obvious to

15 one of ordinary skill in the art, given that one of ordinary skill in the art is

16 creative and has common sense under *KSR*, to then go -- since distance is

17 one of the factors, to go to one that's next to it, the other one, because it's

18 about the same distance from the server?

19   MS. HOLMES: Well, as I said, you have to pick and choose what

20 measurements you are going to maintain and how you're going to use them

21 and certainly, what I'm seeing is that for those references that were cited in

22 this case, when you're talking about looking at proximity, it was always

23 looking at the proximity of the endpoints. So the distance from the client to

24 the content server, the proximity considerations, were not between different

25 and separate requesting clients.

1   So I think that that is a different measurement. I also think that many

2   of the proximity measurements that were cited against us were -- is the client

3   and is the content server, are they in the same continent, those types of

4   measurements, so at a very high level, very gross level. So I think that they

5   are very different in what it's doing.

6   I don't see this combination of using both network or performance-

7   based measurements combined with similarly located client. I don't see that

8   combination taught and I don't believe it's obvious in light of what we've

9   seen.

10   JUDGE HOMERE: What we do know is that the prior art -- the prior

11   art disclose -- discloses you have a plurality of clients and then you have

12   proximity information about them, right, and you use that information. You

13   select the client based on which of the clients is closest. So assuming that

14   proximity information equals cost information, and then if you don't know,

15   would the ordinary, skilled artisan automatically go to the -- defer to the next

16   client if the information for the first client is not available; wouldn't that be a

17   logical conclusion for the skilled artisan?

18   MS. HOLMES: I think that when you're talking about -- I believe, if I

19   can restate my understanding of your hypothetical, that you have multiple

20   clients and you know the locations of those clients, you have multiple

21   servers, you know the locations of those servers. That's the type of

22   information that you would generally have available to you, typically static

23   or semi-static information, and so that if you don't have information about a

24   particular client and a particular server, you're going to go out and figure out

25   what -- where they're located.

1   And so the art that I've seen shows that there is a way of identifying a

2   continent or other very large geographical area that is associated with the

3   location of the content server and the clients. But if you don't have that

4   information, you don't go to another, you simply go out and you find the

5   information that you need, which of course, is an alternate way of doing this,

6   is to say, I don't have information, what I'm going to do is find information

7   relative to this particular client that's making the request.

8   JUDGE HOMERE: No, but let's assume for a moment that you can't

9   find information about the particular client, you can't find proximity

10   information about that particular client, wouldn't the logical thing be to go to

11   the next client to see where you can determine the proximity information

12   and then go from there?

13   MS. HOLMES: Is your hypothetical that you already have proximity

14   information between one client and one content server?

15   JUDGE HOMERE: Well, my hypothetical would require that you

16   don't have proximity information about the first client, but you do have

17   proximity information about the other clients. So wouldn't you

18   automatically, logically go to the next client and then take -- pick the one

19   that's the closest to the server in order to --

20   MS. HOLMES: Well, you -- I don't believe so. I believe what you

21   would do, as I understand the art, is that you would figure out the location of

22   the client that you don't have information for. You would figure out what

23   content server is the closest, is the best fit, that you would do that as opposed

24   to going to a different client, that those references don't talk about

25   maintaining that type of information, to use a different client's information,

1   say information on a second client, to assist with the selection for a first

2   client, that you would simply go out and find that information.

3   And so given any type of a situation where if you don't have data

4   between your two endpoints, to me the more obvious step is, we'll go out

5   and get the data that you need, especially in terms of proximity information

6   which is relatively straightforward to obtain.

7   JUDGE HOMERE: Do you have anything else?

8   JUDGE SIU: Expanding on that hypothetical, the prior art reference,

9   for example, *Colby*, actually doesn't only do proximity; they also have a list

10 of different things that they look at like, I think it says QO -- QOS

11 requirements, degree of load on available servers, network congestion

12 information, and proximity.

13   So given that hypothetical, if you for some reason didn't have

14 information for one client, and you had it for the other ones, given that

15 proximity to the server is one of the features, wouldn't somebody naturally

16 pick another client that you have information for that's about the same

17 distance from the server, meaning close to the client in question that you

18 don't have the information for, because they're about the same distance? So

19 you would know well they're about the same distance, at least I have that --

20 that going for me, and then you could then go and calculate which server is

21 most appropriate? Wouldn't that be kind of a natural thing, to go to the one

22 next to it or very close to it because the distance from the server is the same?

23   MS. HOLMES: I think the distance could be the same, but the

24 direction could be very different, and so just in a -- if you think about the

25 topology of a network, if they're not going through common nodes, you

26 might have distance in terms of -- *Colby* talks about it in terms of a

11

1   continent, so you're somewhat in the same continent, but you may not be

2   using the same path to get from client one to contact server one, client two

3   to content server one. It could be a very different path.

4   So if you're talking about number of hops or something like that, I

5   think that it's not necessarily the same. I understand -- I think in some

6   situations, it might be possible to use that, but it would break down in other

7   situations.

8   JUDGE SIU: I agree it might be different, but the question was,

9   would somebody just choose the one next to it given that it at least has the

10 same distance? I'm not saying that it's going to be exactly the same, the

11 parameters, the cost measurements associated with one and the other. But

12 because we know that one of them is proximity of the client to the server,

13 wouldn't it be a natural thing to go to one that has the same distance, so you

14 have at least that factor the same and/or about the same?

15   The whole idea is, would it have been just natural, you go to the next

16 one given that one of the requirements, the proximity of the client to the

17 server, is the same; I would think it would be kind of a natural choice.

18   MS. HOLMES: I think it also depends on the level of or the type of

19 information you have on location. Again, *Colby* gets it to the continent

20 level, which would put every client in the North American continent likely

21 going to the same content server, and that's exactly what you don't want.

22 You want to be able to disperse the load on these content servers.

23   So I think that given what *Colby* describes, which is simply, I'm going

24 to look for something that's in the same continent, I'm not sure that there's

25 enough information about distance between the client and the content server

26 to be helpful in making an intelligent selection of a content server.

1   JUDGE SIU: Does it only go for something in the same continent? I

2 think if there are multiple ones in the same continent, it would naturally pick

3 the ones closest within the same continent.

4   MS. HOLMES: It may. My recollection of the reference at this time

5 is that it had a table or other indicator that describes continents and that it did

6 that for the purpose of avoiding intercontinental transfers.

7   JUDGE SIU: It assigns a preference of servers which are determined

8 to be closest to the client, follows them, then there's this algorithm. So

9 again, even distance within a continent is taken into consideration, which

10 brings us back to the client-server distance.

11   And given that that's one of the important factors here, how is it not a

12 natural impulse to go to one? If you didn't want to go through the effort of

13 doing this very rigorous calculation, all these other factors, you just want a

14 quick thing, why wouldn't it have been just sort of a natural impulse to go to

15 the next closest client and use that information to apply to the first one?

16 Because proximity and distance is an important thing, and you know, if it's

17 close -- if the second client is close to the first, it's got to be about the same

18 distance to the server.

19   MS. HOLMES: That's correct. I simply don't see any -- any way of

20 getting from -- going -- going out -- if your methodology is to go out and

21 look at each request individually and make decisions on it based on various

22 factors that relate to performance or distance or anything else, I don't see

23 how you go from saying, I'm going to look at each request and I'm going to

24 do these -- evaluate these factors and perform these -- the selection to saying

25 well, I don't have it so I'm not going to do that, but I'm going to do

1   something very different, which still requires some degree of identification

2   of another relevant client or something else.

3   It seems like it is a very different methodology of going out for each

4   particular new client server identification.

5   JUDGE HOMERE: Do you have anything else?

6   MS. HOLMES: I think that there are varying degrees of details that

7   are provided in various claims. We've talked about one of the broader

8   claims today. There are claims that have more discussion of the calculation,

9   the weighted averages and how you determine the distance, also the use of

10   putting in place a closed loop system where once you've made this

11   determination that, okay, I don't have particular information for this client,

12   therefore, I'm going to go to another client that is physically proximate to

13   that client, but now once I go and I make that transaction, I'm going to get

14   data on that transaction, which again relates to network performance and is

15   fed back into this system and it continues on so that it has a closed loop

16   aspect to it. It's not simply okay now, I've picked it and I just have a handful

17   of data points that I'm using on how to select content servers.

18   JUDGE HOMERE: If we turn to claim 32, for example, the

19   distinction you're trying to establish regarding if data is not available, cost

20   measurement is not available, then you go to the next one, that -- different

21   points and it's not recited in claim 32. So I'm at a loss to understanding what

22   is it in claim 32, how would claim 32 be distinguished over the prior art,

23   because that limitation that you're relying on is not even recited in there.

24   MS. HOLMES: Claim 32 relates to the specific calculation that's

25   made between a plurality of clients when you have network performance

26   information for those clients relative to a server, and then there is a

1   determination based on those measurements between those clients and the

2   server, as well as measurements between the various clients, including the

3   client for which you have no information, which is referred to here as the

4   inferable network client, and taking a weighted average of the latency

5   measurements and also considering the distances between the clients to

6   come up with an estimate of latency.

7   So for that inferable network client, you have made a determination as

8   to what the likely latency is between a client for which you have no network

9   performance data and a particular content server.

10   JUDGE HOMERE: I understand that. We just went over what the

11 claim recites, but I don't see -- the argument that you presented for claim

12 one, the same limitation that you just argued from claim one, is not recited in

13 claim 32, for instance. So what is it in claim 32 that would --

14   MS. HOLMES: It is the use of --

15   JUDGE HOMERE: The weighted average?

16   MS. HOLMES: Yes, it is the use of information. In this particular

17 case, it is latency information between clients and a content server and

18 taking that information, as well as information about distance between

19 various clients, to determine what a latency measurement would be for a

20 client where you have no network performance information. So it is related

21 more toward the second element that we discussed with respect to claim one,

22 but it goes into --

23   JUDGE HOMERE : There's nothing in there that discusses not having

24 network performance information. What we have here is we have --

25 inferring -- operational characteristics and then you pretty much provide a

26 series of steps that leads to that. But there's nothing in there that says that

1   what information -- if measurement information is not available at the first

2   client, then you go to the next client and that's how you get it. I don't see

3   anything there that's directed to that.

4   MS. HOLMES: Well, I respectfully submit to you that it is the term

5   "inferable network client" that implies there is no data available for that

6   network client, whereas the other, the plurality of network clients, there is

7   data available for those network clients.

8   JUDGE COURTENAY: But there's no limitation of substituting one

9   client's information for another in your claim; you just have steps of

10   measuring, determining and computing its weighted average?

11   MS. HOLMES: Let's see.

12   JUDGE COURTENAY: I don't see where it's positively recited that

13   you're substituting one client's data for another.

14   MS. HOLMES : I think is, as you say, we don't have a positive step

15   that says we then use that. But it is a method for inferring these

16   characteristics, these operational characteristics for this particular client. So

17   the preamble states that it is a method for inferring operational

18   characteristics associated with the plurality of network clients to an inferable

19   network client.

20   But yes, again to your point, there is no step that would recite what

21   you just pointed out.

22   JUDGE HOMERE: So this claim is incomplete, right?

23   MS. HOLMES: It doesn't include the limitation that has just been

24   raised.

25   JUDGE HOMERE: Therefore, it is incomplete?

26   MS. HOLMES: I believe that the claim calculates --

16

1   JUDGE HOMERE: This -- really, this is a mathematical algorithm,

2   would we not agree, because all you have, you measure something, you

3   determine and you compute. So ultimately what you have is a number.

4   You're calculating a number; that's what this claim is doing?

5   MS. HOLMES: That is certainly part of what this claim is doing and

6   then that number is used --

7   JUDGE HOMERE: Where does it say that it's used -- you measure it,

8   you determine, you calculate, you compute, period?

9   MS. HOLMES: And it's a method for inferring operational

10   characteristics associated with --

11   JUDGE HOMERE: Where is inferring?

12   MS. HOLMES: It's in the preamble.

13   JUDGE COURTENAY: That's really the issue in this claim, as I see

14   it. It's a claim construction issue, how do we construe this term "inferring?"

15   MS. HOLMES: And again it's --

16   JUDGE COURTENAY: How do you think we should construe it? In

17   light of your specification, without reading the limitations from your

18   specification into the claim language, how should we broadly but reasonably

19   construe the term "inferring" in claim 32?

20   MS. HOLMES: And again, it's an issue we didn't have a chance to

21   address directly with the examiner during prosecution. I do see your point.

22   JUDGE COURTENAY: Any other comments you'd like to make,

23   any other issues you'd like to raise?

24   MS. HOLMES: I believe that we've covered the highlights. I know

25   that you've looked at the briefs and the arguments and various things. I think

26   that you have certainly raised some points here today that I wish we'd had an

1   opportunity to talk with the examiner about at different times, but really felt

2   that we were unable to reach agreement on particular things. We've talked

3   about those today, so I believe I've been able to point out to you where we

4   feel that the inventive aspects lie in these claims. So unless you have any

5   other questions, I'm fine.

6   JUDGE HOMERE: Thank you very much.

7   MS. HOLMES: Thank you for your time this morning; I appreciate it.

8   (Whereupon, at 10:47 a.m., the proceedings were concluded.)

9

10